

REMARKS

Reconsideration of the rejections set forth in the Office Action is respectfully requested. By this Amendment, claims 39-40 have been canceled without prejudice or disclaimer and claims 21-22, 26, 29-31, 35, and 38 have been amended. Currently, claims 21-38 are pending in this application.

Election/Restriction

The Examiner required restriction of claims 39-40. Applicants have canceled these claims without prejudice.

Rejections under 35 USC 112

Claims 21-38 were rejected under 35 USC 112, second paragraph. Applicants have amended the claims to overcome this rejection.

Rejections under 35 USC 103

In the Office Action, the Examiner rejected claim 21 under 35 USC 103 as unpatentable over Chen (Flexible Control of a Parallelism in a Multiprocessor PC Router) in view of Applicants Admitted Prior Art (AAPA), and further in view of Venkatanarayanan (U.S. Patent Application Publication No. 2005/0044221), and still further in view of Gourlay (U.S. Patent No. 6,820,123).

This application relates to a method and apparatus for allocating processing capacity of system processing units in an extranet gateway. As discussed by applicants (see e.g. paragraph 7) an Extranet Gateway may be used to connect a VPN site to one or more VPN tunnels. As the number of VPN tunnels supported by a given Extranet Gateway increases, the load on its CPU increases. (See e.g. Paragraph 8). To overcome this issue, Extranet gateways were known to be implemented using multiple CPUs and encryption accelerators (SPUs). (See e.g. Paragraph 9). VPN tunnels would then be assigned to one of the multiple SPUs using a round robin assignment scheme. (See e.g. Paragraph 10).

Applicants recognized that assigning VPN tunnels in a round-robin fashion was not ideal and proposed a more intelligent way of assigning VPN tunnels to System Processing Units (SPUs) in an Extranet Gateway. Specifically, applicants proposed to obtain an initial estimate of

the amount of processing capacity of each SPU and keep track of processing requirements that have been assigned to the SPUs. As new tunnels are required to be established, the tunnel may be assigned by looking to see which processor is estimated to have the most available capacity or the highest relative amount (percentage available) processing capacity.

The combination of references cited by the Examiner does not teach or suggest a method of this nature. Specifically, none of the references teaches or suggests a process in which the initial expected available processing bandwidth of a system processing unit should be used in connection with tunnel assignment, where the initial expected available processing bandwidth represents an amount of VPN tunnel bandwidth that the system processing unit is expected to be able to handle. Applicants describe this concept at paragraphs 21-24. As noted in this section, where the system processing unit is a CPU, the initial expected processing bandwidth for the CPU is based on CPU clock speed multiplied by a conversion factor that has been experimentally or theoretically determined. Where the system processor unit is an accelerator, the expected bandwidth is measured in a test environment and the measured expected processing bandwidth is stored in a table to be conveyed to the load distribution system.

Calculating the initial expected processing bandwidth is more accurate than simply looking at clock speed, since it accounts for how fast the SPU actually is able to process tunnel data. As noted by applicants in paragraph 29, the conversion factor represents an efficiency indication that provides an indication of how efficiently the SPU can handle VPN traffic on a tunnel.

Applicants have amended claim 21 to focus on this aspect. Specifically, applicants have amended claim 21 to recite “establishing a first initial expected available processing bandwidth of a first of the system processing units, the first expected available processing bandwidth representing a first amount of VPN tunnel bandwidth which the first of the system processing units is expected to be able to handle”. Similar amendments have been made to the method step relating to establishing the second initial expected available processing bandwidth of the second system processing unit.

In the rejection, the Examiner cited Chen as teaching CPUs that have different processing speeds which the Examiner contended is related to their throughput. Applicants agree that different CPUs have different processor speeds. However, knowing processor speeds does not tell you how much VPN tunnel bandwidth the system processor unit is likely to be able to

handle. Specifically, as discovered by applicants, since different CPUs may have different conversion factors, similarly clocked CPUs may be able to handle different amounts of VPN tunnel bandwidth. Thus, simply knowing the CPU clock speed does not give an “expected available processing bandwidth representing a first amount of VPN tunnel bandwidth which the first of the system processing units is expected to be able to handle.”

The Examiner cited section 5.2 of Chen as teaching that a CPU can forward 239,234 packets per second which the Examiner contends is related to throughput. In Table 1, Chen describes the cost of forwarding a packet in microseconds. As noted by Chen in section 5.2, it takes 4.18 microseconds of CPU time to forward a packet on a 2 CPU router, from which Chen implies that each CPU should be able to forward 239,234 packets per second.

Applicants have several observations. First, this section is addressing IP processing not tunnel processing. Second, Chen does not use this observation to rate the CPU or create an efficiency factor for the CPU, but rather continues to look to see whether adding additional CPUs would change the amount of data able to be processed by the CPUs. Thus, Chen is not looking to determine how much VPN tunnel bandwidth each processor can handle, but rather is looking to determine whether groups of processors acting together on IP packets can process a larger number of IP packets together as a group than they could individually.

Applicants, by contrast, are processing VPN tunnels by assigning the VPN tunnels to individual SPUs and, to do that, assign an initial expected processing capacity to each SPU. Chen does not teach or suggest a method of this nature. Stated differently, applicants determine the conversion factor for SPU before-hand, and then use this conversion factor to determine an expected processing capacity. Chen does not do this to check the ability of the processor to process packets, but rather is conducting an experiment to see what type of performance can be obtained from the set of processors while running the proposed SMP Click process. Thus, Chen is not testing the performance of the processor, but rather is testing the performance of the SMP Click process.

The deficiencies noted by applicants in connection with Chen are not made up by the other cited references. Certainly, the AAPA does not teach or suggest a method of this nature. Likewise, Venkatanarayan does not appear to teach or suggest anything of this nature.

Gourlay teaches a system that reassigns objects between servers on a network. To do this, Gourlay looks at the capacity of the server in bytes per minute or bytes per second “total

throughput". (Col. 3, lines 5-8). The current throughput (amount of data coming out of the server) is then subtracted from the total throughput to give the available throughput. (Col. 3, lines 9-11). The available throughput is then multiplied by a hotness percentage (e.g. 10%) (Col. 3, lines 12-20) to determine an object threshold value (OTV). Any object that is using more than the OTV is classified as hot, and may be potentially be redistributed to other servers. (Col. 4, lines 55-66). Accordingly, in Gourlay, as the server becomes busier, the available throughput of the server will decrease. This decrease in available throughput will cause the object threshold value to decrease, which will result in a larger number of objects being labeled as "hot" and, hence, a larger chance that some of the objects will be redistributed to other servers.

Gourlay is thus unrelated to this application in that it teaches a way for load to be redistributed between independent servers rather than how to assign the objects in the first instance. More importantly, Gourlay does not teach or suggest that initial expected processing capacity values should be assigned to particular system processing units and used to assign VPN tunnels to the system processing units.

Accordingly, none of the cited references teaches or suggests what is recited in claim 21 as currently amended. Specifically, applicants respectfully submit the references, taken alone or collectively, fail to teach or suggest a method that includes the steps "establishing a first initial expected available processing bandwidth of a first of the system processing units, the first expected available processing bandwidth representing a first amount of VPN tunnel bandwidth which the first of the system processing units is expected to be able to handle" and "establishing a second initial expected available processing bandwidth of a second of the system processing units, the second expected available processing bandwidth representing a second amount of VPN tunnel bandwidth the second of the system processing units is expected to be able to handle". Further, given the teaching of the several references, applicants respectfully submit that it would not have been obvious to use these expected available processing bandwidth estimates in connection with assigning VPN tunnels. Similar amendments have been made to independent claim 30 as well. In view of these amendments, applicants respectfully request that the rejection under 35 USC 103 be withdrawn.

Conclusion

In view of foregoing remarks, it is respectfully submitted that the application is now in

condition for allowance and an action to this effect is respectfully requested. If there are any questions or concerns regarding the amendments or these remarks, the Examiner is requested to telephone the undersigned at the telephone number listed below.

Extension of Time

Applicants hereby request a two month extension of time to respond to the outstanding Office Action. Payment for the extension of time is being submitted concurrently herewith. If any fees are due in connection with this filing, the Commissioner is hereby authorized to charge payment of the fees associated with this communication or credit any overpayment to Deposit Account No. 141315 (Ref: 16263BAUS01U).

Respectfully Submitted

Dated: June 2, 2009

/John C. Gorecki/
John C. Gorecki, Reg. No. 38,471

Anderson Gorecki & Manaras LLP
P.O. Box 553
Carlisle, MA 01741
Tel: (978) 264-4001
Fax: (978) 264-9119
john@gorecki.us